arXiv:1610.06633v1 [cs.HC] 21 Oct 2016

Novelty Learning via Collaborative Proximity Filtering

Arun Kumar
Dept. of Computer Science
University of Minnesota Twin Cities
kumar250@umn.edu

Paul Schrater
Depts. of Psychology and Computer Science
University of Minnesota Twin Cities
schrater@umn.edu

ABSTRACT

The vast majority of recommender systems model preferences as static or slowly changing due to observable user experience. However, spontaneous changes in user preferences are ubiquitous in many domains like media consumption and key factors that drive changes in preferences are not directly observable. These latent sources of preference change pose new challenges. When systems do not track and adapt to users' tastes, users lose confidence and trust, increasing the risk of user churn. We meet these challenges by developing a model of novelty preferences that learns and tracks latent user tastes. We combine three innovations: a new measure of item similarity based on patterns of consumption co-occurrence; model for *spontaneous* changes in preferences; and a learning agent that tracks each user's dynamic preferences and learns individualized policies for variety. The resulting framework adaptively provides users with novelty tailored to their preferences for change per se.

Author Keywords

Recommender Systems; User Behaviors; Boredom; Novelty; User Preferences; Implicit Preferences; Latent Tastes

INTRODUCTION

An exponential growth in online retail and content delivery businesses like Amazon, Spotify, Netflix etc. provided an influx of choices to consumers. The amount of available information at our disposal clearly exceeds the search and consumption capacity of an individual. It creates two conflicting problems - an overabundance of choices but also a saturation of frequent yet not preferred content. In such situations, recommender systems play a crucial role in filtering suitable content for the user. However, recommender systems normally treat preferences as being fixed, an assumption profoundly violated in key recommender domains. Rather than fixed, a user's affinity towards content can change with consumption and time, repetitive or stale content leads to satiation, boredom or devaluation[15], and users are also intrinsically inclined to explore[14].

As a result, in the majority of the current recommender systems, after an initial period of satisfaction, users become increasingly bored of stale and repeated content that do not match their tastes. Dissatisfaction and distrust with the system are the consequence, because recommendations act similarly to *advice* – they are ignored or are greeted with hostility if the system's recommendations show a poor understanding of our tastes and cannot be trusted. Recommenders need to be more like a close aide who knows one's preferences well, whose suggestions are sensitive to our current tastes and to changes in our preferences. For example, phrases like "I am bored of my play list", are common in online music listening. Users start exploring on their own, indicating they have lost trust in the recommender system.

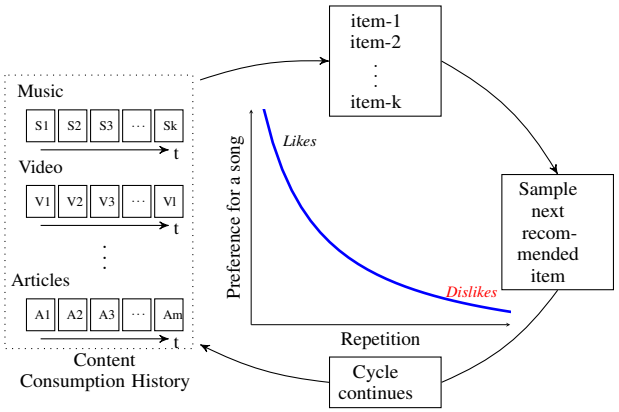

Figure 1. Illustration of the problems spontaneous preference changes create for recommender systems. Under repeated use, systems generate ranked itemized recommendations through a cycle by analyzing consumption histories of users in the system, then updates histories after user consumption. Sooner or later, the recommended item sets become repetitive, and increasingly risk becoming stale and boring to users, due to their spontaneous devaluation of repeated content, shown in the inset graph.

Figure 1 illustrates the problem. Recommender systems generate suggestions by analyzing the collection of consumption histories. The longer users interact with the system, the more likely recommendations will be dominated by a few items that have been frequently consumed, collapsing diversity and novelty in the system. Because items typical devalue with repeated exposure, the system will increasingly generate undesirable recommendations over time. As a result, current systems face two main challenges - tracking the changes in user preferences, and the absence of historical information to recommend new items to a user in the system. Without understanding the current state of users' tastes, it is hard for computational systems to match recommendations to users' liking.

In psychology, changes in human preferences due to boredom or a desire for novelty are well known and characterized[9, 17, 23], but little has been done to incorporate them into recommender systems. Boredom and a desire for novelty have been demonstrated in music listening[15, 14], however, current systems do not have the means to understand or incorporate these latent preferences, largely due to the challenges of estimating

and tracking these cognitive constructs and in using them to improve user's experiences. This creates a fundamental question: how can recommender systems understand when a user is bored or seeking novelty? Another question is how do we determine which user to target for recommendation of new items that do not have any prior history. We answer these two fundamental questions.

To answer these questions on changing preferences with exposure and novel content, we introduce a novel, flexible and modular learning framework to represent user taste, track changes in user taste and learn the individualized policies matching user preferences, termed as *Novelty Learning*. We treat preferences as *dynamically varying consumption bundles*. Items are bundled together so that users prefer items in the bundle to recur in temporal proximity, but we allow the composition of these bundles to change over time. This model of preferences is novel, and we term it *proximity filtering*, because it relies solely on the consumption proximity between items in user's histories. We also find a set of latent musical tastes that are shared across users. We track their dynamic preferences in taste space, which allows us to recommend both novel items to users based on their current taste preferences, and recommend new tastes as the current taste becomes stale. The resulting system is best termed novelty learning via collaborative proximity filtering.

The goal of this paper is to introduce the novel concept of proximity filtering and demonstrate that this method can learn users' individualized novelty preferences in taste space. The remainder of this document presents this system in detail. Some of the common terms used are defined in table I. Section 2 provides background. In section 3 and 4, we discuss the presented model in details. The experimental set up is discussed in section 5, followed by discussion on results in the context in section 6. Finally, we conclude with a note on future scope.

Term	Meaning
Taste	Latent theme discovered from items co-consumed - probabilistic playlists
Taste Profile	User profile based on discovered themes
Preferences	Taste bundles which are dynamic, hence dynamic preferences
Novelty Seeking	User preferences when a user consumes items with changing taste profile
Familiarity Content	User preferences when a user consumes items without a change in taste profile
Proximity Filtering	Similarity by proximity in consumption bundles
NoveltyLearn	Novelty learning agent
Session	A time window of user's music listening activity

Table 1. Terminology

BACKGROUND

With most of the media being consumed in digital form now, online suppliers must target content to customer's preferences. Characterizing user preferences in domains like music that have vast and diverse content poses significant challenges. The sheer diversity requires aggregating content using similarities.

Current recommender systems are driven by inducing similarity: either between users [20], deriving recommendations from other users with the most similar consumption history; or similarity between items. For music recommendations there have been many attempts to exploit both user and item similarity to predict preferences. Oord et.al. [22] propose a latent factor model using weighted matrix factorization and deep convolutional neural networks. They show that recommendations using additional attributes are more sensible and useful when usage data is sparse. Fabian et.al.[2] evaluate tag based user profiles from social media in a personalization effort. Their work focuses on cross-system user modeling to improve recommendation quality. These methods are powerful, but as illustrated in the left panel of Figure 2, this induced similarity relies critically on *static* preferences, while users' preferential tastes are inherently more dynamic .

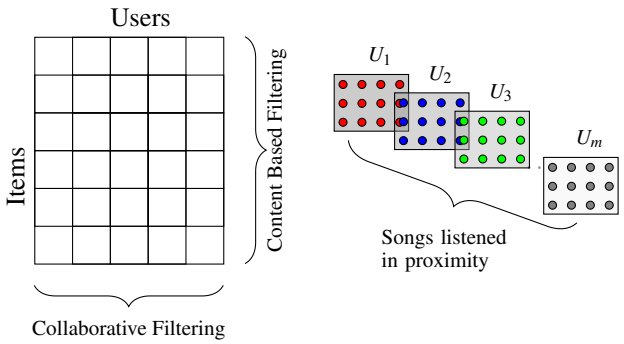

Figure 2. ***Left Panel*** **Collaborative filtering and content based filtering both rely solely on data in user/item matrices that have no explicit representation of time. User/item pairs are usually sparsely sampled and thus change slowly as new user/item choices enter the system. Attribute similarities in content-based filtering are normally predefined and fixed. The result is that both methods implicitly assume near static preferences, creating difficulties in tracking more rapidly varying dynamic preferences.** ***Right Panel*** **Collaborative Proximity Filtering instead relies on finding shared proximity structure across user's song histories. It can be viewed as finding an implicit User/Taste matrix, where tastes are defined by sets of songs listened in close proximity (latent playlists). Recent song choices can be used to track and forecast preference dynamics in taste space.**

Moore et.al.[18] analyze listening preferences of a population of users to try to track a hypothesized gradual change in group preferences over time, ignoring individual tastes. To personalize recommendations, Yin et.al.[24] explore social scenarios and study the influence of other users' choices and ratings on user interests. The rating mechanism, though explicit, is problematic because it is either a delayed reflection of user's state of mind when item was consumed or it can become influenced by other user's choices and public attention. This work also assumes that users' intrinsic interests do not vary much and that temporal dynamics comes from changes in public attention.

In domains like music, users' preferences clearly change individually with time and and should be treated as intrinsically dynamic. This means that all interest can be lost in items that were previously favorites (over-exposure breeds contempt[10] [12]). Users also have well-characterized desire for novel content [3], for which the recommender system will have little or no data. Rapid preference changes like these can catastrophically degrade the user's experience with the recommender system, unless it adapts and tracks these changes. Moreover, different users have different preferences dynamics, and our recommendations also should be adapted to their individualized appetites to explore new options.

In music listening, users have a natural preference dynamics that has previously been identified [15]. Preference for items increases with initial exposure, while boredom (item devaluation) sets in with subsequent exposures. Preferences tend to partially recover with time away from an item. Users recurrently consume a small, working set of items over days and weeks, slowly removing old songs and introducing new ones (essentially an evolving playlist). Viewing a window of listening history as a bundle of items co-consumed, we can say that listening preferences dynamically shift between consumption bundles of songs we term *tastes*. Similarity between songs is induced by being part of the same consumption bundle: similarity via temporal proximity. To better track changes in user's preferences, we exploit this recurrent structure in music listening, using song proximity to filter collaborative and novel recommendations to users (see Figure 2, right panel).

PROXIMITY FILTERING BY TASTES DISCOVERY

Users consume media items like songs in temporal proximity of consumption bundles and not in isolation. These consumption bundles are meaningful in understanding user's preferences at an abstract level. The items within the same consumption bundle are likely to be similar in user's tastes and they are likely to be co-preferred because of the proximity in the consumption bundles as illustrated in Figure 3. It forms the basis of our intuition of proximity filtering.

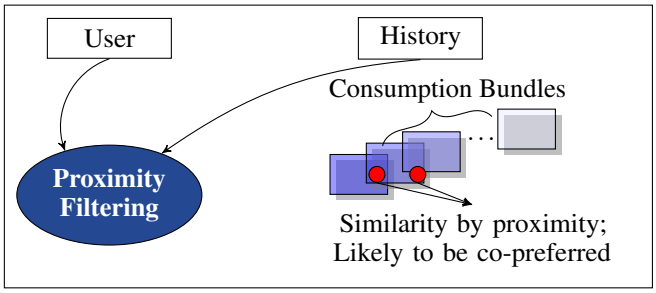

Figure 3. Proximity filtering - items like songs consumed in proximity are likely to be co-preferred

We introduce a novel computational proximity filtering model that considers user tastes as latent structure over dynamic consumption bundles. To discover the underlying common pool of users' tastes, we adapt topic discovery models used in text analysis to discover the equivalent musical topics from listening histories. Topic modeling methods discover useful themes from word co-occurrence frequencies within documents, by

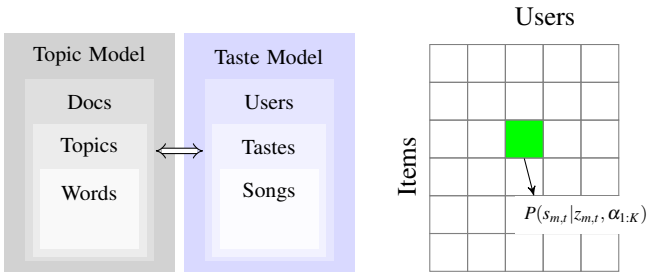

Figure 4. *Left panel* Analogy with standard Latent Dirichlet Allocation: Taste Model treats users content consumption as documents. The discovered taste themes are analogous to topics in text documents and consumed songs are analogous to words. *Right panel* Users consume items/songs according to their tastes. Each cell represents probability of a song in user's taste profile.

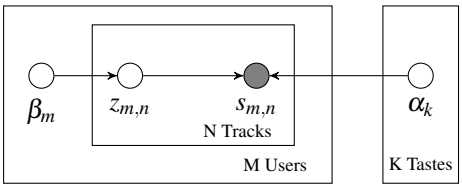

Figure 5. Tastes Discovery: In user music listening history, the songs $s_{m,n}$ are observed. The discovered tastes are common to all the users and lie outside the users' activity space, in plate K tastes. Every user has a proportion z of the discovered tastes. Uncovering these latent tastes is key to the taste space representation.

analyzing a large text corpus. We adapt Latent Dirichlet Allocation [4], a powerful topic modeling technique through an analogy between the set of music listening histories and a text corpus. Figure 4 illustrates this analogy between tastes discovery and a document topic model. Documents correspond to contiguous sections of listening history from a user. The tastes correspond to latent topics in a standard topic model. All the users share the same set of discovered tastes but each user has a probabilistic distribution over these latent tastes. User preferences are dynamic, therefore, we represent a user's *dynamic preference state* as a time-varying mixture of latent tastes. User behavior in content consumption domain is a manifestation of these dynamic preferences.

The discovered taste themes are not predefined but inherently data driven and provide a meaningful and interpretable representation of users' liking. In natural language processing, language modelling methods learn relationships between language constructs. Since our aim is to uncover latent themes to understand user behaviors and liking, topic modeling is an appropriate and preferred method.

More specifically, our tastes model maps into a topic model as illustrated in Figure 5. Each user U_m consumes content $s_{m,n}$ and we can represent the observed content as derived from latent tastes. Each of the user tastes is one among the common pool of the discovered tastes and derived from a taste proportion. Figure 4 illustrates a matrix form with each cell representing the probability of a song in user's taste proportion

and the taste, in turn, is derived from the discovered common pool of latent tastes. Similar to [4], each cell in this matrix represents

$$P(s_{m,n}|z_{m,n},\alpha_{1:K}) = \alpha_{z_{m,n},s_{m,n}} \quad (1)$$

Discovered tastes as well as taste proportions for users are assumed to come from a Dirichlet. Therefore, the model is reduced to finding a joint distribution similar to [4]

$$P(\alpha_{1:K},\beta_{1:M},z_{1:M},s_{1:M}) = \Pi_{i=1}^{K} P(\alpha_i)\Pi_{m=1}^{M} P(\beta_m) \\ (\Pi_{n=1}^{N} P(z_{m,n}|\beta_m)P(s_{m,n}|z_{m,n},\alpha_{1:K})) \quad (2)$$

Until recently, users were largely profiled using static attributes like gender, location etc., but such static attributes are not of much value [1]. We uncover user attributes by understanding the distinctive mixture of their latent tastes as well as the similarities in a common taste pool. This interpretation lies at the core of taste space transformation to learn *inherent* user preferences and novelty seeking behaviors.

Our model consists of modular building blocks providing flexibility, re-usability and extensibility to a range of content delivery domains. We present the schematic diagram of the framework in Figure 6. In the taste discovery, we obtain tastes distribution for users and tracks distribution for tastes. The tracks distribution for tastes demonstrate that different tastes vary in proportions of tracks. As we observe, the tastes have differentiating proportions of a set of tracks and these proportional compositions characterize the discovered tastes. While the discovered tastes are shared across users, users have a distinctive proportional mixture of latent tastes. This representation allows us to transform users' itemized consumption into latent taste space using the discovered tastes pool and learn users' novelty seeking behaviors. In a nutshell, the discovered latent tastes are leveraged to learn user preferences and we demonstrate our method for users' online music listening.

NOVELTY LEARNING

In Psychology, boredom is often considered as an emotion of dullness or lacking stimulation. Goetz et.al. [11] describe and support the external validity of boredom with varying degrees of arousal and valence. However, due to its hidden presentation in perceived environment, it is hard to track boredom using explicitly observable signals like ratings of items, satisfaction scores etc. In an academic setting, Nett et al. [19] explored different strategies with groups of individuals to find out their activities in boredom situations. In absence of any visual or observable cue about user's emotional state, recommender systems fail to understand whether a user is novelty seeking or not and users, in absence of content matching their taste, get bored sooner than later.

Figure 7 illustrates devaluation of user preferences as user receives recommendations. In the beginning, users might have a stronger liking to the recommended items, however, as the cycle of repetitive or stale recommendations continues, users get bored and stop trusting the system leading to either user churn or disengagement. Boredom increases with repeated exposure. Since boredom does not manifest itself as visibly as other emotional states like anger or excitement but a more silent one, novelty seeking users might not even provide a

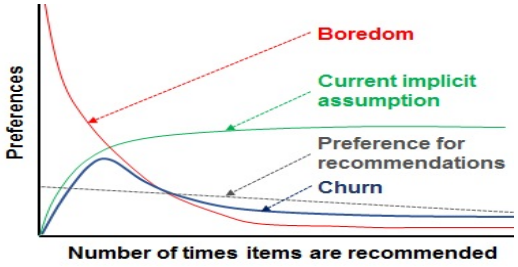

Figure 7. Devaluation of user preferences for recommended items implies that the boredom settles in and user might not trust the system any longer. From systems' point of view, without any proactive intervention to improve user satisfaction, this user might stop using the provided service altogether. Therefore, a proactive mechanism to learn which users are seeking novelty proactively is critical.

feedback or any directly observable signal. It makes current recommender systems completely clueless to understand users' tastes. Once a user is disengaged, it is an uphill task for the systems to earn user's trust back. Therefore, we need a machinery to understand user behaviors more proactively and we aim to bridge this fundamental gap by providing a method to learn novelty seeking behavior using latent tastes.

Taste Membership

Users consume contents at items level like songs but their implicit behaviors and preferences can not be understood in item level view but an abstract view. The items are consumed in bundles which reflect users' affinity to a taste. Such a representation is useful in understanding implicit constructs that drive preferences.Therefore, we use the discovered tastes pool to associate each of the item bundles with a taste and transform temporal sequences of users' consumed items bundles into an abstracted taste space. Figure 8 illustrates the method. We

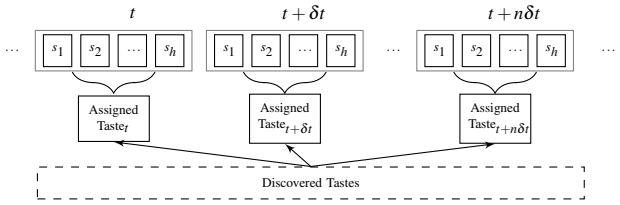

Figure 8. Taste Membership Assignment - temporal windows of consumed items are transformed into taste space using discovered tastes pool.

compute taste association i.e. a taste to be assigned to a temporal window. Let s_i represent song or track id, k, tastes and h be the history of tracks listened in a time period of length n given by $[s_1, s_2, ...s_n]$. For every temporal window, we compute the probability that a taste is associated with it and for a history h, it is given by

$$P(k|h) = \Pi_{i=1}^{n} P(s_i|k)P(k)/\Sigma_{i=1}^{n} P(s_i|k)P(k) \quad (3)$$

where $P(k)$ = user taste probability and $P(s_i|k)$ = taste track probability This evaluated probability is used to assign a taste to each temporal window i.e. session as

$$Taste_{assigned} = max_k P(k|h) \quad (4)$$

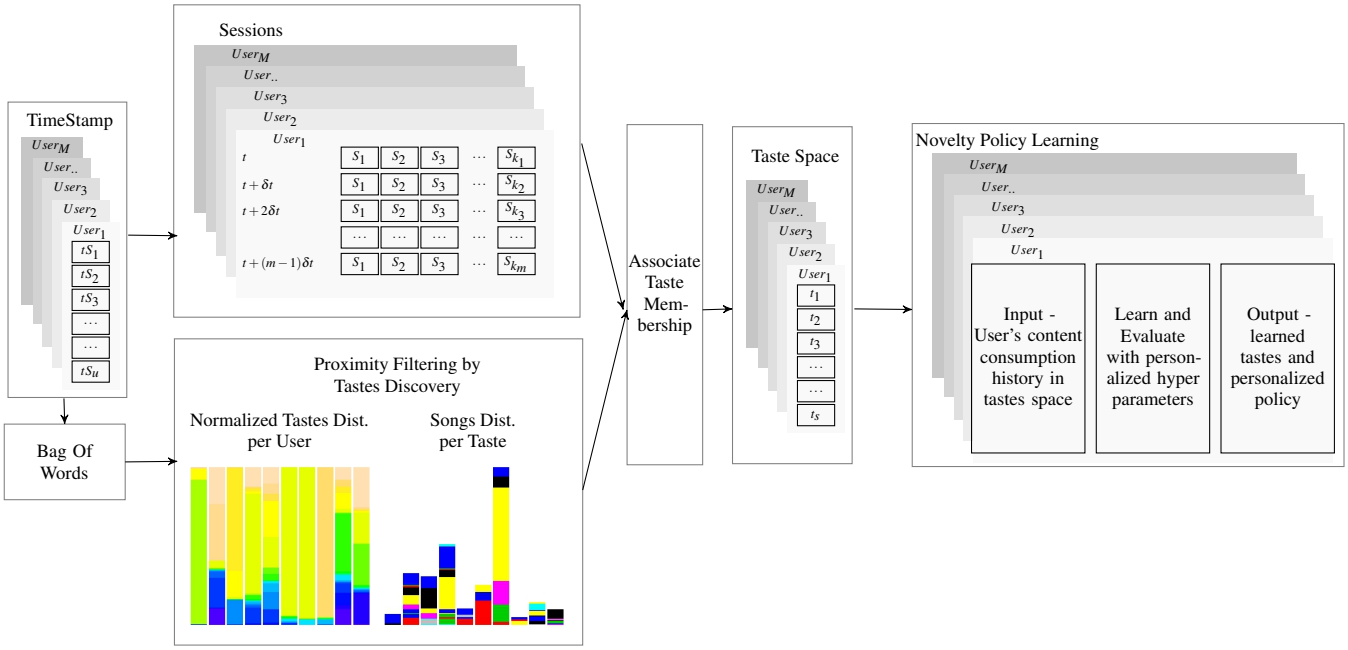
TimeStamp
Sessions
Bag Of Words
Proximity Filtering by Tastes Discovery
Normalized Tastes Dist. per User
Songs Dist. per Taste
Associate Taste Membership
Taste Space
Novelty Policy Learning
Input - User's content consumption history in tastes space
Learn and Evaluate with personalized hyper parameters
Output - learned tastes and personalized policy

Figure 6. Model: Users consume items at a granular level but their liking or preferences are abstract and implicit. We introduce proximity filtering by discovering the latent tastes. The itemized space is transformed into a taste space to provision for learning users' novelty seeking behaviors.

Consequently, users' items consumption space is transformed into a latent taste space and we can now learn users' preferences.

Novelty Policy Q-Learning

Learning users' behaviors like novelty seeking is crucial but ignored aspect in not only recommender systems but also human factors and engagement initiatives. An implicit nature of boredom manifestation also introduces difficulty in computational tracking of user behavior. To answer this problem, we aim to learn the users' novelty seeking policies. Since item level consumption view does not present insights abstract enough to learn behaviors, we need to resort to a taste space representation. Users tastes are dynamic and we present changing tastes as novelty seeking behavior, it also represents boredom with content. When users continue to consume items without a change in their tastes, they are content with the familiarity. Our novelty learning agent, a model free reinforcement learning agent, learns users' novelty seeking policies.

Reinforcement learning [21] has its roots in psychology and animals' reward seeking behavior. The goal is to maximize the utility by taking a set of actions in an environment state and learn the policy that maximizes the utility. The environment is typically conceptualized as a Markov Decision Process. The reinforcement learning methods differ from classic supervised learning in a way that we do not know the correct input and output pairs nor there is any attempt to make an explicit correction. Rather, the models are based on experiencing the reward and exploitation-exploration trade-offs to find out the optimal or maximizing utility. Such level of unknown input-output pairs make it suitable to our learning problem. We use a model free reinforcement learning method called Q-learning [13] in which we do not need to know any transition probability model

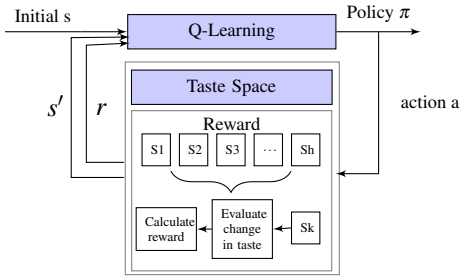

Figure 9. Novelty learning agent is a model-free Q-learning agent in taste space. It learns an optimal policy π^* on whether a user likes to seek novelty or content with familiarity.

beforehand. With Q-learning, for each of the users, we learn state-action mapping that gives the expected utility of taking an action in a given state and in turn, it gives us an optimal policy. Cohen et.al. [6] show that the exploratory preferences depend on user's state. We learn behaviors from users' taste states. In an episode, an action is taken, a reward is realized and an update to Q values is made. The magnitude of update depends on the learning rate and discount factor. We consider the state, s, at time t as user's tastes representation at time t. As illustrated in Figure 9, our policy learning agent learns an optimal policy in a taste state, s, with actions a, novelty seeking or content with familiarity, taking the user to a new state, s' and it receives a reward. The cumulative reward is maximized, thereby, we learn an optimal policy which is a mapping of preferences with state as illustrated in the update equation below:

$$Q(s,a) := Q(s,a) + \alpha[r + \gamma \max_{a'} Q(s',a') - Q(s,a)] \quad (5)$$

where α is learning rate and γ discount factor. Learning rate determines to what extent the values will be updated and discount factor determines the importance of future rewards. We set α as a decaying function of iterations, γ as 0.9 and reward values are empirically set for users. Once the Q-learning converges, we get an optimal policy in a state, $\pi^*(s)$, given by

$$\pi^*(s) = \arg\max_a Q^*(s,a) \quad (6)$$

and this learned policy reflects users' novelty seeking behaviors. *Finally*, our presented model learns a crucial but hard to learn implicit user preferences like boredom or novelty seeking. Due to its flexible and modular design, the presented novelty learning method could play a significant role in making systems perceptive and trust worthy in a longer term.

Evaluation Metrics

Our measurement metrics are F1Score, Accuracy and Value of Personalization, as defined below:

F1Score and Accuracy

Given that tp is true positive, tn true negative, fp false positive and fn false negative, we can define F1Score and Accuracy as

$$P = tp/(tp+fp); R = tp/(tp+fn) \quad (7)$$

$$F1Score = (2PR)/(P+R) \quad (8)$$

$$Accuracy = (tp+tn)/(tp+tn+fp+fn) \quad (9)$$

Value of Personalization

User behaviors are individualistic and intrinsic which makes personalization of great value. For every user, the proposed model learns an optimal policy for novelty seeking preferences and we propose a metrics to understand the value of such personalization. Our intuition is to know how much a user would benefit going by a self policy than other user's policies. For every user, we capture F1Score using policy of other users other than self. We define this metrics Δ as

$$\Delta_i = Average(F1Score_{ij} - F1Score_{ii}) \forall i \neq j \quad (10)$$

The greater this difference, the better the value of personalization for that particular user.

EXPERIMENT

Our experiment is based on a publicly available users' music listening history dataset provided by last.fm [5]. Last.fm is a popular music listening service helping users explore a variety of music, learn new music, connect with community and revisit previously listened tracks. Since this dataset has a temporal data, it forms a natural choice for our experiment requiring consumption diversity. Dataset [16] consists of complete music listening history of around 1K users from 2005 till 2009 with user id, time stamp, artist id, artist name, track id and track name for every user. The dataset statistics is provided in table 2. Algorithm 1 presents the key steps involved.

Attribute	Records	Unique Users	Artists	Unique Tracks
Count	19,150,868	992	176,948	961,416

Table 2. Dataset Statistics

Algorithm 1 Proximity Filtering and Novelty Learning

```
Parameters: K number of discovered tastes, w temporal window size, r reward values
x ← Input bag-of-words format {pre-processing}
α, β = LDA(x) {get User Tastes and Taste Tracks dist.}
X_{m,t,s} = Input content consumption in temporal windows {pre-processing}
X_{m,t} = Associate Taste Membership
for j = 1 : M do {for each user}
Training
    Initialize RL Parameters
    while not converged do
        select action a with ε-greedy policy from Q
        take action a
        s', r ← new state, reward
        make Q-update {Novelty Policy Q-Learning}
        s ← s'
    return policy
Evaluate Evaluate on held out set
```

In pre-processing and sessions management phases, we convert users' music listening data from a naive format to time series as well as bag-of-words format as required to discover tastes pool. An off-the-shelf package [7] is used for topic modeling. To assess a reasonable set of model parameters with stability, we computed the perplexity of the held out user set. The held-out set was unseen by the model during training. For a trained model, the perplexity of the unseen set is given by:

$$Perplexity(U_{test}) = exp\{-\Sigma_{i=1}^{m} log(p(u_i|model))/\Sigma_{i=1}^{m} N_i\} \quad (11)$$

where $log(p(u_i|model))$ is the log likelihood, calculated by fitting the trained model to the held-out set. We varied number of topics from 1 to 40 and number of iterations for convergence as 10,100,1000 and 10000. We observed that the perplexity becomes stable beyond the number of iterations for convergence set as 100 iterations. We chose the number of topics to 20 which is at the beginning of a flat perplexity region and iterations as 1K because there is no gain in perplexity beyond these values. Moreover, we set up the temporal session window size to one hour empirically. We varied it from 30 minutes - 1 week respectively. A larger window size does not reflect users' 'in-moment' taste because users listen to tracks in proximity depending upon their current mood or state of mind. Using these parameters, we discover tastes to get tastes distribution per user and tracks distribution per taste and associate each of the temporal sessions with a taste. Therefore, users' items consumption history is transformed into a taste space.

The taste space representation is at an abstract level and we use it to learn novelty seeking Q-learning policy. Reinforcement learning requires learning rate and discount rate which are set as $1/(t^{0.65})$ and 0.9 respectively following several insights on convergence [8]. For experimental simulation, the reward values are set in a range from -1 to 1 as well as a small action cost. The Q-learning iterates until convergence and learns an optimal policy, whether a user is going to be novelty seeking or not given the current taste state of that user. We validate

the learned policy using 80-20 training-validation split. Since the users' music listening data is in a time series format, this 80-20 split up is suitable when compared to other cross validation methods. We evaluate F1Score, Accuracy and Value of Personalization for every user.

In a supporting experiment to demonstrate users' quitting behavior and how taste representation can be beneficial, we use taste distribution per user from tastes discovery and create users' taste profiles. This is similar to the previous experiment except that instead of associating a taste, we associate a taste profile. Our aim, now, is to gain some insights into users' listening patterns with their taste profile similarities with previous sessions. Using this similarity score, we demonstrate that learning novelty in latent taste space is insightful in understanding users' content consumption behavior and knowing the potential customers who can quit.

To the best of our knowledge, there is no other computational method to learn inherent user tastes and dynamic behaviors like boredom or novelty seeking which are not directly observable. Therefore, a direct comparison with any method would be incomprehensible. However, to check our model, we compare our policy learning measures with that of Support Vector Machine with rbf kernel. Finally, we evaluate the value of personalization quantitatively.

RESULTS AND DISCUSSION

We evaluated our model on the last.fm music listening dataset. We discuss user churn followed by novelty learning.

User Quit and Novelty Seeking Behavior in Taste Space

User churn is a major concern in a number of domains. Proactive knowledge of users who seek novelty is an important step to keep their trust in the recommender systems. We look at users' quitting behavior with respect to their taste profile sequences. We compute an average taste profile similarity score by comparing users' taste profile in the current session with that in the previous session. This measure is used to demonstrate our taste model's ability to learn user behavior using two example users - one who quit and stopped using the services within 449 days and other who continued to consume music listening services. Figure 10a demonstrates that the average number of tracks listened in a month by the 'quitting' user continued to decline and the system was unable to act on this behavior because it had no clue if the user is seeking novelty or not. As shown in Figure 10b, the average profile similarity for this 'quitting' user is lower with $\mu = 0.51$ and $\sigma = 0.15$ compared to the 'continuing' user with $\mu = 0.89$ and $\sigma = 0.11$. It demonstrates that our novelty learning method in *taste space* could play a pivotal role in understanding user behaviors proactively.

We also sought to understand an approximate user quiting behavior at a subgroup level. To achieve that, users were categorized based on their activities time period. We considered quitting users as the users whose activity time period from their first recorded activity to the last recorded activity was less than a year and their last recorded activity was at least one year prior to the end of the dataset time period. This ensured that we do not incorrectly categorize the users who started

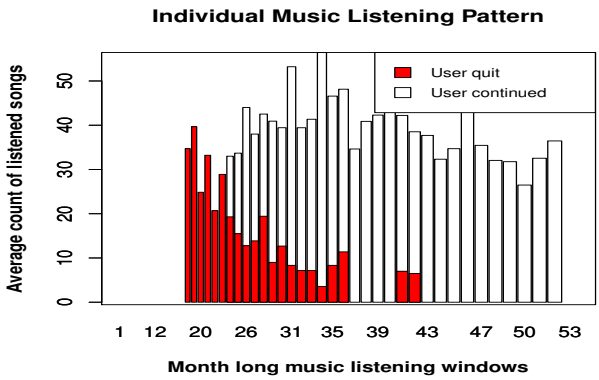

a. Songs listened in monthly windows

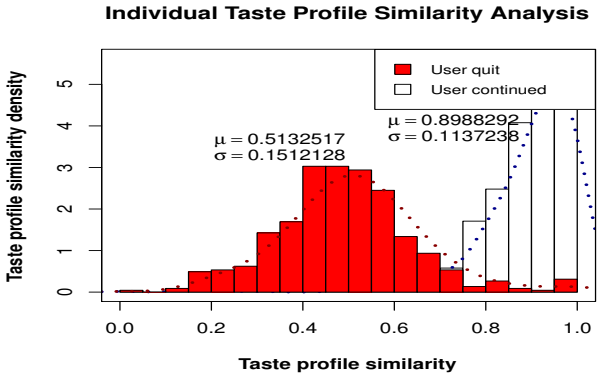

b. Taste profile similarity varies with the type of user

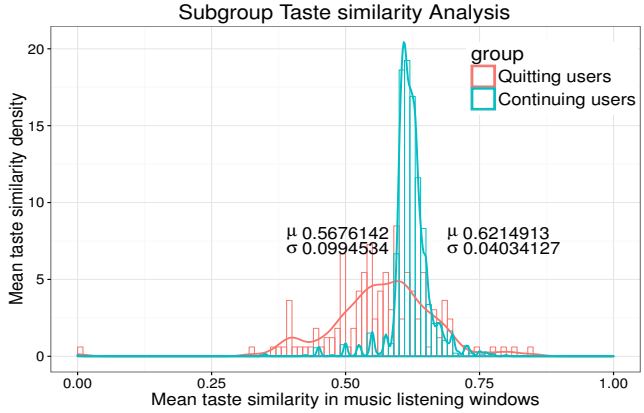

c. Two subgroups have different taste similarities

Figure 10. Taste space representation is adequate to proactively know which users are seeking novelty. *a* demonstrates that the quitting user's listening activity continued to decline in comparison to the continuing user. *b* compares taste profile similarity scores for the two users and demonstrates that quitting user was seeking more novelty. *c* is subgroup level taste profile similarities comparison of quitting and continuing users. It demonstrates our ability to learn novelty in taste space.

using the service very late. Similar to the individual user behavior analysis, we computed the taste similarity scores of users but at an aggregate level and demonstrate the subgroup level measures in Figure 10c. As evident, the quitting users has a lower mean taste profile similarity score implying a more

exploratory behavior for this subgroup. It further establishes that the taste space representation is adequate and abstract enough to understand user's novelty seeking behavior.

Novelty Policy Q-Learning

Taste space representation of consumed content is used to learn users' novelty seeking behavior. Indirectly, learning novelty seeking behavior in taste space also means learning users' boredom with currently delivered content. Having such a representation and learning method is useful in content delivery systems involving human behaviors. With this view, we train our novelty learning agent, NoveltyLearn. We keep 20% of temporal data for validation and evaluate two measurement metrics F1Score and Accuracy. The average evaluation

↓Model Metric→	F1Score	Accuracy
NoveltyLearn	0.80326	0.87763
SVM	0.77700	0.83429

Table 3. Novelty Learning agent learns an optimal policy whether a user is seeking novelty or content with familiar items.

F1Score of novelty seeking policy learner is 0.80326. Due to our model's ability to learn implicit tastes, any comparison with existing methods is incomprehensible but we compare average F1Score and Accuracy using single taste state representation of our method with a popular machine learning technique Support Vector Machines with RBF kernel using a longer taste history as input features. As shown in table 3, our learning agent, despite using only a single taste state with rewards, does better than SVM on an average by 2.6 points on F1Score and 4.3 points on Accuracy measure. The policy learner converges to an optimal policy, see Figure 11.

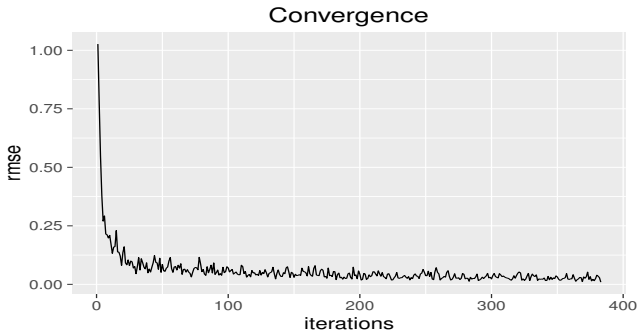

Figure 11. Novelty learning agent converges to an optimal policy.

Evidently, the presented model is capable of learning user behavior in latent tastes space. The learned policy is useful to proactively understand user behavior and take action before they loose trust in system and quit. The novelty policy learning agent makes no consideration of user profile attributes like age, gender, location etc. which is a *significant step* in establishing user profiles purely based on their consumption behavior.

Value of Personalization

The value of personalization measure demonstrates how beneficial it is to consider user personalization. To evaluate it, we consider policies learned for users and apply them across other users except self on a 80-20 validation set. It means that for a user i, we evaluate the learner's policies according to all other users except i to know how does it perform when it is not using its own policy. From this matrix form, the value of personalization is derived using the definition in equation

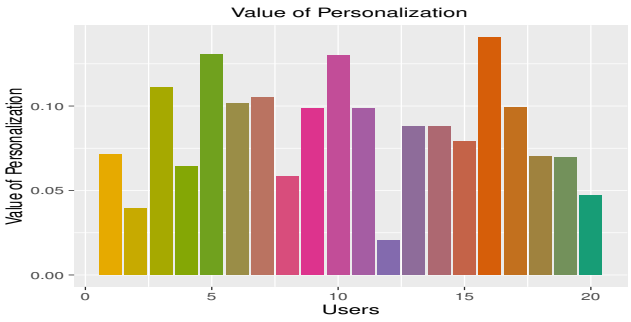

Figure 12. Value Of Personalization demonstrates that personalization is beneficial.

10. As shown in Figure 11, an overall average gain for a user using own policy is around 8 points. We observe that one user has a lower gain and it is attributed to a globally set parameter, length of sessions, to maintain evaluation consistency across users. Therefore, we state that it is beneficial to pursue personalization.

Market Benefit

Current recommender systems consider users' behavioral preferences as fixed and it is hard for computational systems to track unobservable behaviors like boredom. With an ability to learn such implicit behaviors, our method has a huge potential in a number of domains like music, movies etc. that rely heavily on user tastes and diversity of choices.

CONCLUSION

Our method fills a fundamental gap in learning implicit user behaviors. After an initial period of interaction with the system, users get bored with stale and repetitive content that does not match their implicit tastes and lose trust in the system. Therefore, an ability to learn users' tedium state is significant for user retention and satisfaction.

We presented a novel method called *collaborative proximity filtering* by leveraging patterns of consumption co-occurrence. The richness in understanding users' implicit and latent states that is provided by the taste model is not feasible in items space. We presented a learning agent that learns *novelty* seeking behavior in taste space. A proactive knowledge of such a user behavior is vital so that systems can take actions before a user actually churns and leaves the system. Our method provides a new way to profile users based on their taste learnt from the consumption proximity rather than the static attributes. Further, we demonstrated that personalization is useful and should not be ignored in recommender systems. The presented model brings complex psychological constructs within the purview of computation and a novel paradigm in creating a synergy between implicit behaviors and computational methods by learning from the data.

REFERENCES

1. 2016. Netflix's Grand, Daring, Maybe Crazy Plan to Conquer the World. (2016). http://www.wired.com/2016/03/netflixs-grand-maybe-crazy-plan-conquer-world/
2. Fabian Abel, Eelco Herder, Geert-Jan Houben, Nicola Henze, and Daniel Krause. 2013. Cross-system user modeling and personalization on the social web. *User Modeling and User-Adapted Interaction* 23, 2-3 (2013), 169–209.
3. D.E. Berlyne. 1960. *Conflict, Arousal and Curiosity*. McGraw-Hill.
4. David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *the Journal of machine Learning research* 3 (2003), 993–1022.
5. O. Celma. 2010. *Music Recommendation and Discovery in the Long Tail*. Springer.
6. Jonathan D Cohen, Samuel M McClure, and Angela J Yu. 2007. Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 362, 1481 (2007), 933–942. DOI: http://dx.doi.org/10.1098/rstb.2007.2098
7. Lda developers. 2015. LDA Python. (2015). https://pypi.python.org/pypi/lda
8. Eyal Even-Dar and Yishay Mansour. 2004. Learning rates for Q-learning. *The Journal of Machine Learning Research* 5 (2004), 1–25.
9. P James Geiwitz. 1966. Structure of boredom. *Journal of Personality and Social Psychology* 3, 5 (1966), 592.
10. Moshe Givon. 1984. Variety seeking through brand switching. *Marketing Science* 3, 1 (1984), 1–22.
11. Thomas Goetz, Anne C. Frenzel, Nathan C. Hall, Ulrike E. Nett, Reinhard Pekrun, and Anastasiya A. Lipnevich. 2013. Types of boredom: An experience sampling approach. *Motivation and Emotion* 38, 3 (2013), 401–419. DOI: http://dx.doi.org/10.1007/s11031-013-9385-y
12. Abel P Jeuland. 1979. Brand choice inertia as one aspect of the notion of brand loyalty. *Management Science* 25, 7 (1979), 671–682.
13. Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. 1996. Reinforcement learning: A survey. *Journal of artificial intelligence research* (1996), 237–285.
14. Komal Kapoor, Vikas Kumar, Loren G. Terveen, Joseph A. Konstan, and Paul R. Schrater. 2015. "I like to explore sometimes": Adapting to Dynamic User Novelty Preferences.. In *RecSys*, Hannes Werthner, Markus Zanker, Jennifer Golbeck, and Giovanni Semeraro (Eds.). ACM, 19–26.
15. Komal Kapoor, Nisheeth Srivastava, Jaideep Srivastava, and Paul Schrater. 2013. Measuring spontaneous devaluations in user preferences. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1061–1069.
16. Last.fm. 2015. last.fm. (2015). http://www.last.fm
17. Edgar Pessemier Leigh McAlister. 1982. Variety Seeking Behavior: An Interdisciplinary Review. *Journal of Consumer Research* 9, 3 (1982), 311–322. http://www.jstor.org/stable/2488626
18. Joshua L Moore, Shuo Chen, Douglas Turnbull, and Thorsten Joachims. 2013. Taste Over Time: The Temporal Dynamics of User Preferences.. In *ISMIR*. Citeseer, 401–406.
19. Ulrike E. Nett, Thomas Goetz, and Lia M. Daniels. 2010. What to do when feeling bored?: Students' strategies for coping with boredom. *Learning and Individual Differences* 20, 6 (2010), 626 – 638. DOI: http://dx.doi.org/10.1016/j.lindif.2010.09.004
20. Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*. ACM, 285–295.
21. R.S. Sutton and A.G. Barto. 1998. *Reinforcement learning: An introduction*. Vol. 116. Cambridge Univ Press.
22. Aaron van den Oord, Sander Dieleman, and Benjamin Schrauwen. 2013. Deep content-based music recommendation. In *Advances in Neural Information Processing Systems 26*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 2643–2651. http://papers.nips.cc/paper/5004-deep-content-based-music-recommendation.pdf
23. Stephen J Vodanovich. 2003. Psychometric measures of boredom: A review of the literature. *The Journal of psychology* 137, 6 (2003), 569–595.
24. Hongzhi Yin, Bin Cui, Ling Chen, Zhiting Hu, and Zi Huang. 2014. A temporal context-aware model for user behavior modeling in social media systems. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*. ACM, 1543–1554.